On the Excess Bandwidth Allocation in ISP Traffic Control for Shared Access Networks

Kyeong Soo Kim, *Member, IEEE*

arXiv:1403.6272v1 [cs.NI] 25 Mar 2014

***Abstract*—Current practice of shaping subscriber traffic based on token bucket by Internet service provider (ISP) allows short-term fluctuations in its shaped rate and thereby enables a subscriber to transmit traffic at a higher rate than a negotiated long-term average. The traffic shaping, however, results in significant waste of network resources, especially when there are only a few active subscribers, because it cannot allocate excess bandwidth to active subscribers in the long term. In this letter we investigate the long-term aspect of resource sharing in ISP traffic control for shared access networks. We discuss major requirements for the excess bandwidth allocation in shared access networks and propose ISP traffic control schemes based on core-stateless fair queueing (CSFQ) and token bucket meters. Simulation results demonstrate that the proposed schemes allocate excess bandwidth among active subscribers in a fair and efficient way, while not compromising the service contracts specified by token bucket for conformant subscribers.**

***Index Terms*—Access, Internet service provider (ISP), traffic shaping, fair queueing, quality of service (QoS).**

I. Introduction

THE practice of shaping subscriber traffic by Internet service provider (ISP) has been under intensive study; for example, the effect of ISP traffic shaping on various packet-level [1], [2] and user-perceived [3] performances has been investigated, which provides a new insight into the actual performance of broadband access networks.

One critical issue is that traffic shaping cannot allocate excess bandwidth to active subscribers in the long term. This is because the traffic shaper based on token bucket cannot take into account the status of other subscribers. As extensively studied in [1], [3], a large-size token bucket enables sharing of excess bandwidth among active subscribers, but only in the short period of time corresponding to the token bucket size.

The modification of token bucket algorithm to allocate excess bandwidth has been studied in the context of fair queueing/scheduling [4], [5] and differentiated services (DiffServ) networks [6]. The results of these studies, however, cannot be applicable to the current ISP traffic control which is not based on DiffServ. Also, the modification of token bucket algorithm and/or the change of its negotiated parameters during the operation may raise the issue of traffic conformance — which is currently based on the original token bucket algorithm — and compromise the quality of service (QoS) of conformant traffic as a result.

A desirable alternative to the traffic shaping based on a modified or adaptive token bucket would be the use of the original token bucket as a meter in order to separate traffic from a subscriber into conformant and non-conformant one and treat them differently in further processing based on the status of a network and other subscribers (e.g., [7], [8]). The issue of per-subscriber allocation of excess bandwidth proportional to its negotiated long-term average rate, however, has not been studied in this context.

In this letter we discuss major requirements for the excess bandwidth allocation in shared access networks and propose ISP traffic control schemes based on core-stateless fair queueing (CSFQ) [9] and token bucket meters that can meet the requirements.

K. S. Kim is with the College of Engineering, Swansea University, Swansea, SA2 8PP, Wales United Kingdom (e-mail: k.s.kim@swansea.ac.uk).

II. Excess Bandwidth Allocation

A. Requirements

We define the excess bandwidth in downstream at time t for an access network with N subscribers as follows:

$$C_{ex}(t) \triangleq C - r_c(t), \tag{1}$$

where C is the capacity of the access link and $r_c(t)$ is the arrival rate of conformant packets for all the subscribers from the network.[1] Below we set two major requirements that any excess bandwidth allocation schemes should meet:

- The allocation of excess bandwidth should not compromise the QoS of subscribers' traffic conformant to service contracts based on the original token bucket algorithm.
- Excess bandwidth should be allocated among active subscribers proportional to their negotiated long-term average rates, i.e., token generation rates.

The first requirement is more fundamental than the second one because both subscribers and ISPs consider the excess bandwidth allocation as an optional feature and therefore its benefit should not come at the expense of other subscribers; note that the traffic conformance is solely based on the ISP traffic control at the edge of the network and covers access links only. The second requirement, on the other hand, enables ISPs to provide new service and pricing schemes with more incentives to subscribers willing to pay more for higher long-term average rates.

B. ISP Traffic Control Schemes based on WFQ and CSFQ

Fig. 1 shows an access switch for a shared access network. Considering the requirements in Sec. II-A, one can come up

[1]The discussions in this letter are also applicable to upstream traffic with minor modifications because the upstream traffic control in shared access networks is also centralized and located in the access switch (e.g., using grants in cable Internet and Ethernet passive optical network (EPON)).

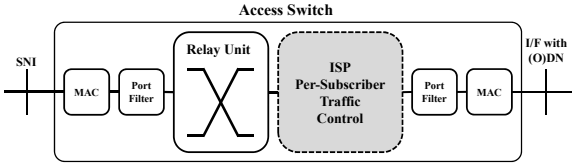

Fig. 1. Block diagram of an access switch for a shared access network.

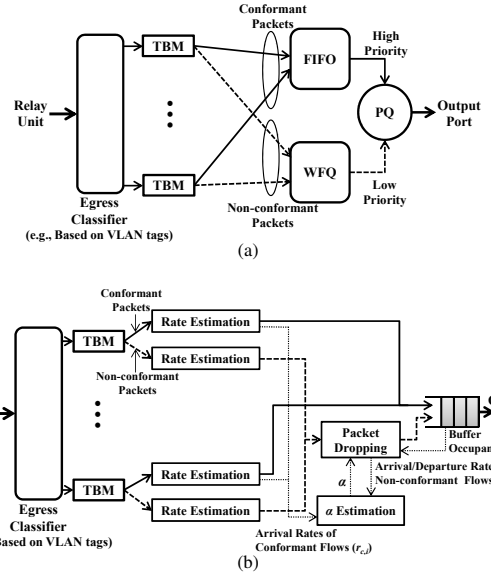

Fig. 2. ISP per-subscriber traffic control enabling proportional allocation of excess bandwidth: (a) A conceptual model based on WFQ and (b) a practical implementation based on CSFQ.

with a conceptual model of ISP per-subscriber traffic control shown in Fig. 2 (a), which enables proportional allocation of excess bandwidth based on weighted fair queueing (WFQ) and priority queueing (PQ) with token bucket meters (TBMs): The first requirement is met by the use of token bucket meters and PQ with higher priority for conformant packets, while the second requirement is met by WFQ. Note that, even with PQ ahead, WFQ can still maintain its fairness property [10].

This conceptual model based on WFQ, however, has a major flaw: Due to the separation of traffic from the same subscriber into two flows and separate queueing, packet sequence is not preserved, which makes it impractical for user datagram protocol (UDP) applications. Fig. 2 (b) shows a practical implementation based on CSFQ, which can preserve packet sequence through a common first in, first out (FIFO) queue. The architecture shown in Fig. 2 (b) corresponds to the extreme case of CSFQ islands, i.e., the node itself is an island. Because both edge and core router functionalities reside in the same node, there is no need to carrying labels in packets between the rate estimation and the packet dropping units.

Let $A(t)$ be the total arrival rate of non-conformant packets at time t, i.e., $A(t) \triangleq \sum_{i=1}^{N} r_{nc,i}(t)$, where $r_{nc,i}(t)$ is the arrival rate of non-conformant packets for the ith subscriber. If $A(t) > C_{ex}(t)$, the *normalized* fair rate $\alpha(t)$ is a unique solution to

$$C_{ex}(t) = \sum_{i=1}^{N} w_i \min(\alpha(t),\ r_{nc,i}(t)/w_i), \tag{2}$$

On receiving *packet* p *for the* i*th subscriber*
if `Conform(`p`)` $== True$ **then**
  $r_c \leftarrow$ `EstimateRate(`r_c`, Length(`p`))`
  `Enque(`p`)`
  $C_{ex} \leftarrow$ `ExcessBW(`r_c`)` `/* using (1) */`
else
  $r_{nc,i} \leftarrow$ `EstimateRate(`$r_{nc,i}$`, Length(`p`))`
  $prob \leftarrow$ `Max(`$0,\ 1-\alpha(w_i/r_{nc,i})$`)`
  if $prob >$ `UniformRandom(`*0, 1*`)` **then**
    $\alpha \leftarrow$ `Estimate`α`(`$r_{nc,i}/w_i$`,Length(`p`),`***True***`)`
    `Drop(`p`)`
  else
    $\alpha \leftarrow$ `Estimate`α`(`$r_{nc,i}/w_i$`,Length(`p`),`***False***`)`
    `Enque(`p`)`

Algorithm 1: Pseudocode of rate estimation and packet dropping.

where w_i is the weight for the ith subscriber, which is proportional to the token generation rate; otherwise, $\alpha(t)$ is set to $max_i\,(r_{nc,i}(t)/w_i)$ [9]. Based on the excess bandwidth, arrival rates, and normalized fair rate, we can now implement rate estimation and packet dropping as described in Algorithm 1, which is a modified version of weighted CSFQ with two arrival rates per subscriber: If $r_{nc,i}(t)/w_i \le \alpha(t)$, a non-conformant packet will be enqueued for forwarding; otherwise, the packet will be dropped with the probability of $\max(0,\ 1-\alpha(w_i/r_i))$.

The estimation of the normalized fair rate (i.e., $\hat{\alpha}$ for α) is described in Algorithm 2, where $\hat{A}$ and $\hat{F}$ are the estimated aggregate arrival rate and the estimated aggregate rate of the accepted traffic of non-conformant packets, respectively, and K_α is a window size to filter out the inaccuracies in rate estimation. The update of the estimator $\hat{\alpha}$ is based on linear approximation of the function $F(\cdot)$, i.e., $\hat{\alpha_{new}} = \hat{\alpha_{old}} \times C_{ex}/\hat{F}$.

As discussed in [9], we use exponential averaging to estimate various rates, i.e., r_c, $r_{nc,i}$, A and F, whose general formula is given by $x_{new} = \left(1 - e^{-T/K}\right)\frac{l}{T} + e^{-T/K}x_{old}$, where T is elapsed time since the last update, which means the interarrival time of corresponding packet, and K is an averaging constant (K_α for A and F).

To better support bursty, elastic traffic like that of transmission control protocol (TCP), we can also implement buffer-based amendment as in [9]. When receiving a packet, we check the buffer level against a predefined threshold. Every time the buffer level passes the threshold, we decrease $\hat{\alpha}$ by a small percentage (9% for the simulation in this letter). Note that the major purpose of this amendment in the current scheme is to prevent non-conformant traffic from hogging the buffer space of the common FIFO queue at the expense of conformant traffic, unlike that of the original CSFQ.

III. Simulation Results

We carried out a comparison study of the proposed scheme with the conceptual model as a reference. For WFQ implementation in the conceptual model, we use deficit round-

Function Estimateα(*r, l, dropped*)
Data: r is the normalized arrival rate, l a packet length, and $dropped$ a flag indicating whether the packet is dropped or not. $\hat{\alpha}$ and r_{max} are initialized to C_{ex} and 0, respectively.
Result: $\hat{\alpha}$ (i.e., fair share rate) is returned.
$\hat{A} \leftarrow$ EstimateRate($\hat{A}$, l)
if $dropped == False$ **then**
  $\hat{F} \leftarrow$ EstimateRate($\hat{F}$, l)
if $\hat{A} \geq C_{ex}$ **then**
  if $congested == False$ **then**
    $congested \leftarrow True$
    $start_time \leftarrow current_time$
    if $\hat{\alpha} == 0$ **then** $\hat{\alpha} \leftarrow$ Min(r, C_{ex})
  else
    if $currentt_time > start_time + K_\alpha$ **then**
      $\hat{\alpha} \leftarrow \hat{\alpha} C_{ex}/\hat{F}$
      $start_time \leftarrow current_time$
      if $\hat{\alpha} == 0$ **then** $\hat{\alpha} \leftarrow$ Min(r, C_{ex})
else /* $\hat{A} < C_{ex}$ */
  if $congested == True$ **then**
    $congested \leftarrow False$
    $start_time \leftarrow current_time$
    $r_{max} \leftarrow 0$
  else
    if $currentt_time < start_time + K_\alpha$ **then**
      $r_{max} \leftarrow$ Max(r_{max}, r)
    else
      $start_time \leftarrow current_time$
      $\hat{\alpha} \leftarrow r_{max}$, $r_{max} \leftarrow 0$
return $\hat{\alpha}$

Algorithm 2: Pseudocode of fair rate estimation.

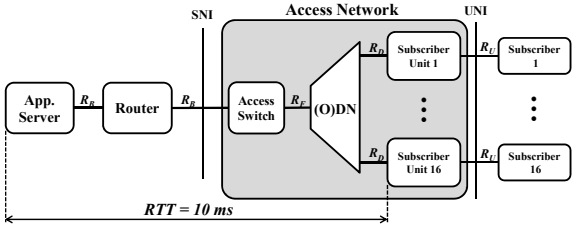

Fig. 3. A simulation model for a shared access network with 16 subscribers.

robin (DRR) [11]. Fig. 3 shows a simulation model where 16 subscribers are connected through 100-Mb/s user-network interfaces (UNIs) to shared access with the same feeder and distribution rates of 100-Mb/s, each of which receives packet streams from UDP or TCP sources in the application server. The backbone rate (i.e., R_B) and the end-to-end round-trip time are set to 10 Gb/s and 10 ms. To model the shared (optical) distribution network ((O)DN), an Ethernet switch with the same feeder and distribution rates is used because feeder and distribution links are identical (e.g., cable Internet) or passively connected in a star topology (e.g., EPON) in shared access. The implementation details are given in [3].

We divide 16 subscribers into 4 groups (i.e., 4 subscribers per group): For Groups 1-3, each subscriber receives a 1000-byte packet at every 0.5 ms (i.e., the source rate of 16 Mb/s) from a UDP source. Token generation rate, however, is set to 2.5 Mb/s for Group 1, 5 Mb/s for Group 2 and 7.5 Mb/s for Group 3. We also set starting time to 0 s, 60 s, 120 s, respectively. For Group 4, each subscriber receives packets from a greedy TCP source with token generation rate of 10 Mb/s and starting time of 180 s. Token bucket size is set to 1 MB for all subscribers, and peak rate control is not used at all. The size of FIFO and per-subscriber queues of DRR is set to 1 MB (i.e., 17 MB in total) for the reference scheme (denoted as "DRR+TBM"), and the size of common FIFO queue is set to 16 MB for the CSFQ-based scheme without ("CSFQ1+TBM") and with buffer-based amendment ("CSFQ2+TBM") to cope with worst-case bursts resulting from 16 token buckets with size of 1 MB each; as for the buffer-based amendment, we set a threshold to 64 kB. The averaging constants used in the estimation of flow rates (i.e, K) and the normalized fair rate (i.e., K_α) are set to 100 ms and 200 ms, respectively.

Fig. 4 shows flow throughput averaged over a 1-s interval from one sample run, which demonstrates dynamic performances of each scheme (i.e., *how quickly* it can respond to the changes in incoming traffic and allocate excess bandwidth accordingly). Until 180 s when TCP flows start, all three schemes can allocate available bandwidth (including excess bandwidth) among UDP flows well, with DRR+TBM being the best in terms of fluctuation and convergence speed. Due to 1-MB token buckets, there are spikes in the throughput of newly started flows at 60 s (i.e., Group 2) and 120 s (i.e., Group 3), while the throughput of existing flows temporarily plunged accordingly. As TCP flows start at 180 s, the difference among the three schemes become clearer: Because packet sequence is not preserved in DRR+TBM, which causes lots of retransmissions, throughput of TCP flows fluctuate most. With CSFQ1+TBM, while the fluctuation in TCP flow throughput is not so big, the convergence is quite slow (about 10 s to reach the token generation rate of 10 Mb/s). In this regard we found that the buffer-based amendment in CSFQ2+TBM efficiently reduces the transient period, especially for TCP flows, at the slight expense of fluctuations in steady states.

Fig. 5 shows the average throughput of flows for two 50-s periods (i.e., a subperiod (60 s) minus a transient period (10 s)) with 95 percent confidence intervals from 10 repetitions, demonstrating static performances of each scheme (i.e., *how exactly* it can allocate available bandwidth among subscribers per the requirements described in Sec. II-A in a steady state). As shown in Fig. 5 (a), both DRR+TBM and CSFQ1+TBM allocate excess bandwidth from Group 4 exactly per (2), while CSFQ2+TBM suffers from the fluctuations observed in Fig. 4 (c). With TCP flows, however, CSFQ2+TBM performs best and guarantees well the negotiated long-term average rates for newly started TCP flows, even though the difference among the schemes is not that big. Note that dotted lines indicate the fair share of each flow.

IV. Conclusions

In this letter we have studied the long-term aspect of resource sharing in ISP traffic control for shared access networks

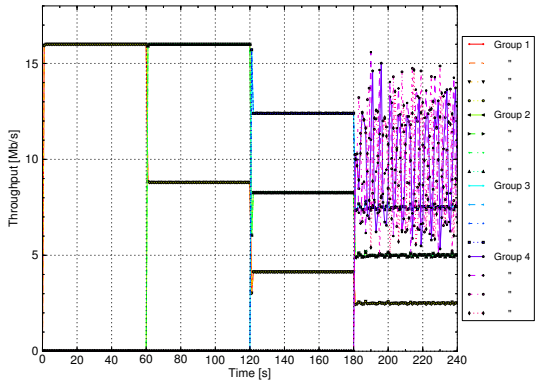

(a)

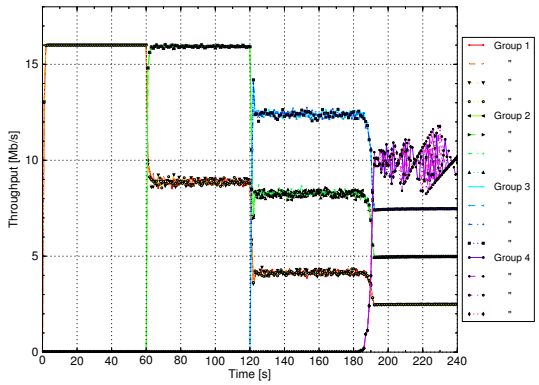

(b)

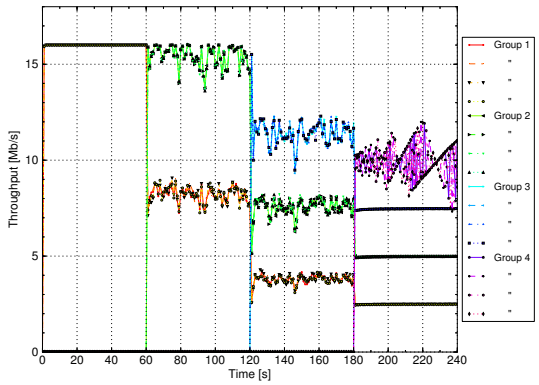

(c)

Fig. 4. Time series of throughput of flows: (a) DRR+TBM, (b) CSFQ1+TBM, and (c) CSFQ2+TBM.

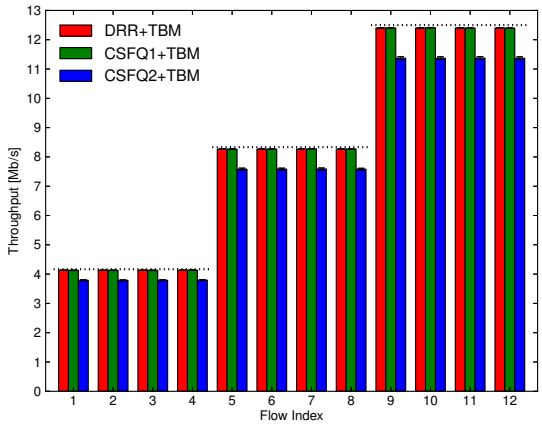

(a)

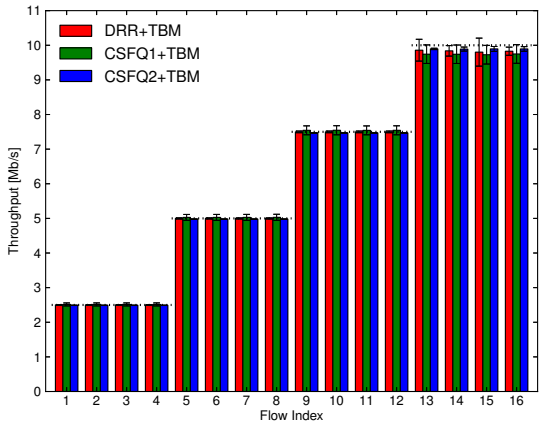

(b)

Fig. 5. Average throughput of flows with 95 percent confidence intervals for the period of (a) $130 \leq t < 180$ and (b) $190 \leq t < 240$.

and proposed ISP traffic control schemes based on CSFQ and token bucket meters. Simulation results demonstrate that the proposed schemes allocate excess bandwidth among active subscribers in a fair and efficient way, while not compromising the service contracts specified by the token bucket algorithm for conformant subscribers. With buffer-based amendment, we could reduce transient period of the proposed scheme and thereby improve throughput of interactive TCP flows.

REFERENCES

[1] S. Bauer, D. Clark, and W. Lehr, "PowerBoost," in *Proc. HomeNets'11*. New York, NY, USA: ACM, Aug. 2011, pp. 7–12.

[2] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè, "Broadband Internet performance: A view from the gateway," in *Proc. SIGCOMM'11*, Toronto, Ontario, Canada, Aug. 2011, pp. 134–145.

[3] K. S. Kim, "The effect of ISP traffic shaping on user-perceived performances in broadband access networks," in *Proc. ICUMT 2012*, Petersburg, Russia, Oct. 2012, pp. 533–538.

[4] D. Abendroth, M. E. Eckel, and U. Killat, "Solving the trade-off between fairness and throughput: Token bucket and leaky bucket-based weighted fair queueing schedulers," *AEU - International Journal of Electronics and Communications*, vol. 60, no. 5, pp. 404–407, 2006.

[5] J. Kidambi, D. Ghosal, and B. Mukherjee, "Dynamic token bucket (DTB): A fair bandwidth allocation algorithm for high-speed networks," in *Proc. 1999 ICCCN*, Boston, MA, USA, Oct. 1999, pp. 24–29.

[6] E.-C. Park and C.-H. Choi, "Adaptive token bucket algorithm for fair bandwidth allocation in DiffServ networks," in *Proc. 2003 IEEE GLOBECOM*, vol. 6, San Francisco, CA, USA, Dec. 2003, pp. 3176–3180.

[7] Y. Huang, R. Guérin, and P. Gupta, "Supporting excess real-time traffic with active drop queue," *IEEE/ACM Trans. Netw.*, vol. 14, no. 5, pp. 965–977, Oct. 2006.

[8] B. Patt-Shamir, G. Scalosub, and Y. Shavitt, "Competitive analysis of buffer policies with SLA commitments," in *Proc. ICNP 2008*, Orlando, FL, USA, Oct. 2008, pp. 197–206.

[9] I. Stoica, S. Shenker, and H. Zhang, "Core-stateless fair queueing: A scalable architecture to approximate fair bandwidth allocations in high-speed networks," *IEEE/ACM Trans. Netw.*, vol. 11, no. 1, pp. 33–46, Feb. 2003.

[10] Y.-C. Wang and Y.-C. Tseng, "Packet fair queuing algorithms for wireless networks," in *Design and Analysis of Wireless Networks*, ser. Wireless Networks and Mobile Computing, Y. Pan and Y. Xiao, Eds. Nova Science Publishers, Inc., 2005, vol. 1, ch. 7, pp. 113–128.

[11] M. Shreedhar and G. Varghese, “Efficient fair queueing using deficit round-robin,” *IEEE/ACM Trans. Netw.*, vol. 4, no. 3, pp. 375–385, Jun. 1996.